arXiv:1610.05985v1 [cs.CV] 19 Oct 2016

Robust Video Synchronization using Unsupervised Deep Learning

Ido Freeman* Patrick Wieschollek* Hendrik P.A. Lensch
University of Tübingen

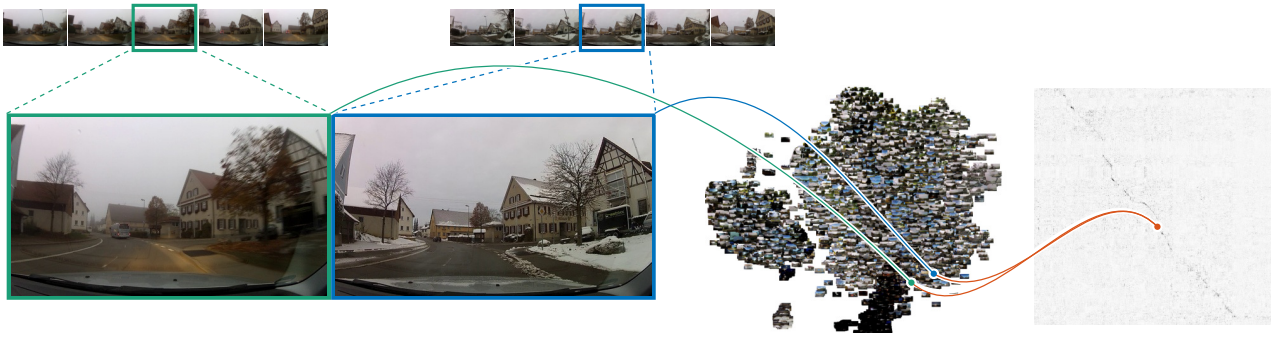

Figure 1: *Our novel unsupervised training algorithm is used to train a convnet to temporally align two videos captured in different environment settings. The neural network learns a content-aware similarity metric by embedding each frame into a feature space which is independent from factors like weather, traffic or seasons. For the extraction of a video alignment, we deploy a Dijkstra-based algorithm.*

Abstract

Aligning video sequences is a fundamental yet still unsolved component for a wide range of applications in computer graphics and vision. Especially when targeting video clips containing an extensively varying appearance. Using recent advances in deep learning, we present a scalable and robust method for computing optimal non-linear temporal video alignments. The presented algorithm learns to retrieve and match similar video frames from input sequences without any human interaction or additional annotations in an *unsupervised* fashion. An iterative scheme is presented which leverages on the nature of the videos themselves in order to remove the need for labels. We incorporate a variation of Dijkstra's shortest-path algorithm for extracting meaningful training examples as well as a robust video alignment. While previous methods assume similar settings as weather conditions, season and illumination, our approach is able to robustly align videos regardless of such noise. This provides new ways of compositing non-seasonal video clips from data recorded months apart.

Keywords: Deep Learning, Video Editing, Alignment, Synchronisation

1 Introduction

Temporal *video alignment* (synchronization) enables various applications such as, but not limited to, video stitching for creating panoramic videos [Agarwala et al. 2005], merging multiple takes of a scene into a single video [Rüegg et al. 2013], automatic summarisation of videos [Ngo et al. 2005] and HDR video generation [Kang et al. 2003]. These algorithms are heavily dependent on a robust and accurate temporal alignment of video-frames between multiple videos. Aligning videos means to find a dense frame-to-frame mapping for each time step in the involved videos with a consistency over time.

One straightforward way to align videos is by manually annotating sparse landmarks in frames using a video editing software. The video sequences are usually represented as "draggable" bars on a global time-line. The relative position of these bars represents a temporal relation between the clips. Yet, using such a tool to create markers and drag video clips requires extensive human interaction which could be time-consuming and error-prone. In addition, these tools make a strong, not necessarily true assumption, namely that the content between two landmarks is homogeneous and linear along the time dimension. This means that users are limited to specifying the start and end points of video clips. Then, given a set of landmark positions, the synchronization step then usually boils down to a linear interpolation with an accuracy dependent on the density of annotated frames.

One way to automate the process would be to record a synchronization signal during acquisition. Proposed solutions for this meta-information are Longitudinal Time Codes, genlock, GPS data or a landmark-based audio fingerprinting [Bryan et al. 2012]. Given this kind of temporal mapping, it becomes possible to *spatially* match the content from multiple frames using local descriptors. Solutions based on GPS data are as accurate as the very device which registered the data. Although these methods seem to produce acceptable to good results, in many real-world applications these explicit synchronisation signals are not available as most consumer cameras only encode the creation date of the video file in the meta-data.

In our setting, we use a collection of videos containing mostly overlapping driving routes over the course of a year under different daytime illuminations, varying of traffic conditions, i.e. temporary road works, rush hour congestions, etc., diverse weather conditions, i.e. clear skies, snowfall, rain, and seasonal environment appearance, i.e. effects of vegetation. Additionally, the videos mostly show journeys between the same two cities but still have a variation in start and end locations and the actual roads driven. Videos are captured by a dashboard camera without a fixed spatial mounting or any further synchronization meta-information like key-point frames. The temporal alignment of the videos thus poses a further challenge when trying to match stationary situations like waiting in traffic during rush hour with a video showing little traffic. The extensive variation in the data thus requires a non-linear algorithm like Dynamic Time Warping [Wang et al. 2014] to *stretch a video clip* to be in sync with its counterpart. Indeed, using a collection of

*Indicates equal contribution.

common methods like Transfer Learning and exploiting the temporal coherence in videos [Mobahi et al. 2009], our approach learns to handle arbitrary videos.

The enormous range of possible variations in video content which also suffers from a noisy acquisition (e.g., different viewing angles, wipers, raindrops, camera transformation, etc.) requires context understanding of the entire scene rather than a simple local feature matching like histogram based methods using SIFT features [Wang et al. 2014]. Moreover, the videos might show interrupted content, e.g. when the lens is cleaned or the camera is remounted during recording. To robustly synchronize our videos, we propose a novel approach which is based on learning a content-based similarity metric by a Deep Convolutional Neural Network [LeCun et al. 1995] in an unsupervised manner, i.e. without any sort of annotation.

The presented model consists of three supporting steps.

- The similarity ranking of a given frame pair.
- A robust way of filtering those similarities such that undesired, structured noise is removed.
- A globally optimal path extraction from the filtered data.

With these steps a high precision training data is iteratively and automatically harvested from the video data improving with each iteration while learning to cope with the increasing complexity of the data.

The main contribution of this paper are:

– An *unsupervised* learning-based method for matching video frames from different videos under varying situations without annotations or synchronizations signals.
– A direct application of similarity-learning using *Artificial Neural Networks* to replace an exhaustive nearest neighbour search over high dimensional feature vectors.
– A robust computation of a matching tour for partially overlapping video pairs and automatically detected start and end points of the matching tour.

2 Related Work

The process of video alignment holds a natural relation to image alignment which was addressed by several studies, e.g. using stereo correspondence estimation [Scharstein and Szeliski 2002] or robust pixel descriptors [Liu et al. 2011].

Yet in contrast to the image alignment problem, for matching video frames several additional challenges as moving content, changes in perspective, illumination arise. While minor alignment errors would not be visible in vanilla image alignment, they become notable when replaying both videos synchronously.

Basic video alignment is commonly used in the field of human action retrieval or surveillance motion capture. Hereby, finding similarities of human actions in videos are based on dynamic time warping of various sensor features to track the human skeleton [Zhou and De la Torre Frade 2009].

This work is not the first to address spatio-temporal video alignment. [Lei and Yang 2006; Wolf and Zomet 2006; Tuytelaars and Gool 2004] have published a their take on tackling the issue. However, the proposed algorithms all assume a linear temporal correspondence.

[Caspi and Irani 2002; Ukrainitz and Irani 2006] have looked into multi-sensor Alignment of motion and actions in videos. Yet, this, in a rather similar manner to [Lei and Yang 2006; Wolf and Zomet 2006; Tuytelaars and Gool 2004], assumes a constant time shift between successive frames and does not address static scenes.

Figure 2: *The 100-th frame from 9 random videos in our dataset, covering the high diversity of different lightning and weather conditions from dusk, bright sunshine to night-time and snow. Notice, the blurred image caused by rain drops.*

Sand and Teller [2004] compute a matching-likelihood describing the 3d motion of a physical 3d point, which takes several minutes to compute for a second of video content. This enormous amount of computational complexity renders large-scale video matching infeasible.

A non-linear flexible solution was proposed by Wang et al. [2014]. In their paper, they compute a matching histogram of SIFT features using approximate nearest neighbour search. Yet, for the synchronisation of two 500 frames long videos from our dataset, their method required a nearest neighbour search over 10^7 vectors in a 128 dimensional feature space which is computationally expensive. Comparing to [Wang et al. 2014], our approach represents single a frame using a 1024-dimensional vector, allowing for a very fast direct computation of similarity matrices and content-based video search.

3 Learning to Synchronize Videos

In this section, we shortly discuss our dataset and then continue with describing our proposed method to synchronize a video pair. We look at a pair of videos a, b from a collection $\mathcal{V}$ of videos ($a, b \in \mathcal{V}$). For the sake of simplicity, we only describe the synchronization of video pairs. The further extension of this method to multiple videos is then rather straightforward. For a synchronization of each pair of videos (a, b) we assume a pairwise partial overlap., i.e. video a should contain some frames, which are similar to frames of video b.

3.1 Dataset

The underlying dataset comprises of 1204 full-HD, 30FPS videos (1.8 TB of raw data) capturing a commuters' car journeys (each takes roughly 100min) captured on partially overlapping routes in Germany between April 2012 and March 2013 and spanning over approximately $16,000$km. The videos were captured using a GoPro Hero 3 camera mounted on the dashboard before every journey.

The dataset features both rural landscapes and urban scenes under various weather conditions all throughout the mentioned time period. Figure 2 illustrates randomly sampled video content. We have not applied video rectification, stabilization or any other kind of

pre-processing in our approach except for down-scaling and mean subtraction.

3.2 Algorithm Outline

In this section, we shortly outline the learning approach to motivate the specific steps before providing more details regarding the actual implementation. The algorithm consists of two phases: the *interpretation* phase and the *discovery* phase.

Interpretation-Phase Given a frame-pair (a_t, b_s) from videos a, b, one could either regard it as similar ($y_{ts} = 1$) or as different ($y_{ts} = 0$). In section 3.5 we describe how to automatically infer the labels y_{ts} for each frame-pair. For now we assume that y_{ts} is already given. We train a convolutional neural network to create a high dimensional vector representation of each frame, such that the *matching distance* space between two arbitrary frames is small for (a_t, b_s) if and only if the given frame-pair is similar, i.e. $y_{ts} = 1$, and outside of a certain radius for different frames. This gives a learned matching distance

$$\delta \colon \mathcal{V} \times \mathcal{V} \to \mathbb{R}_+, \quad (a_t, b_s) \mapsto \delta(f(a_t), f(b_s)),$$

In our case δ is the Euclidean metric $\| \cdot \|_2$ of the content-representations $f(a_t), f(b_s)$.

Discovery-Phase This phase *automatically harvests* new training data for the interpretation phase. We *only* assume the matching-distance δ to be known. Now, for each pair of videos (a, b) with T_a and T_b frames, respectively, a cost matrix $C \in \mathbb{R}_+^{T_a \times T_b}$ of non-negative entries can be computed by

$$c_{ij} = \delta(f(a_i), f(b_j)), \quad i = 1, \dots, T_a, \; j = 1, \dots, T_b,$$

representing the matching distances of all frame-pairs in the videos. See the top row of Figure 5 for some examples of the cost matrix C obtained from this formulation.

Using a modified version of the Dijkstra algorithm and additional heuristics which are discussed in subsection 3.4, we get a temporal alignment of the videos. Given feedback from the temporal alignment we could now reliably detect and label matching video frames. Namely, the frame pairs along the matching tour are considered as positive training data $(a_s, b_t, y_{st} = 1)$ for fine-tuning the neural network.

Alternating both phases (i.e. *interpretation* and *discovery*) in an iterative fashion yields a robust algorithm for aligning seasonal videos without any human intervention.

Throughout the rest of this section, we explain these steps in more detail.

3.3 Model Design

Prior to describing the architecture of our neural network, we shortly recap the idea of deep neural networks and their training process.

Probabilistic models label each possible configuration (frame-pairs) by a specific normalized matching probability. Maximizing the likelihood of a path would give the most probable matching tour. Unfortunately, computing these probabilities contains computationally challenging normalization terms. Instead we train a convolutional neural network [LeCun et al. 1995] to directly compute the a similarity metric $\delta(\cdot, \cdot)$, which can be seen as associating an energy to input pattern pairs. A neural network $f^{(\theta)}(\cdot)$ with a set of

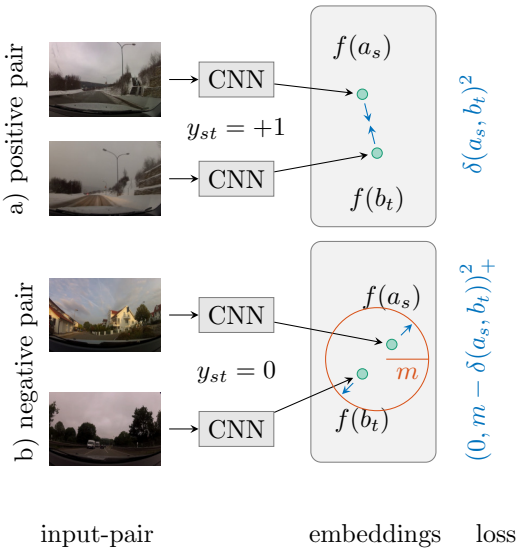

Figure 3: *Let (a_s, b_t) be a frame-pair. During training, the neural network is optimized to similarly embed positive pairs (a) while negative pairs are pulled away from each other, if their distance is smaller than a margin m.*

parameters θ is a composition of parametrized functions

$$f^{(\theta)}(x) := (F_n^{(\theta)} \circ F_{n-1}^{(\theta)} \circ \cdots \circ F_1^{(\theta)})(x).$$

Hereby, each function $F_k^{(\theta)}$ represents a layer in the neural network, e.g. a linear operation $x \mapsto \hat{\theta} x$, a non-linearity $\tanh(\cdot)$, a convolution with learned filters or a max-pooling operation. Some of these layers contain trainable parameters, which are updated depending on the loss from the prediction, which we will cover now.

To evaluate the distance function between two video frames, we set δ as the Euclidean distance between two frames embedded by the network.

$$\delta(a_s, b_t) = \| f^{(\theta)}(a_s) - f^{(\theta)}(b_t) \|_2$$

We first embed each frame of the pair using the same convnet in a Siamese manner [Bromley et al. 1994]. The embedding vectors of both frames are then used to compute the gradients wrt. the contrastive loss function [1] [Hadsell et al. 2006]. We therefore optimize the network parameters θ to be

$$\arg\min_{\theta} \underbrace{y \frac{1}{2} \delta(a_s, b_t)^2}_{\substack{\text{penalty for similar frames} \\ \text{with large embedding distance}}} + \underbrace{(1 - y) \frac{1}{2} (0, m - \delta(a_s, b_t))_+^2}_{\substack{\text{penalty for dissimilar frames} \\ \text{with small embedding distance}}} \tag{1}$$

for some margin $m \in \mathbb{R}$. The first term penalizes positive frame-pairs with embeddings that are too far away from each other. The latter penalizes negative frame-pairs if they are too close to each other (closer than m).

We use the Siamese neural network approach [Bromley et al. 1994] to efficiently train a concept of similarity which generalizes to unseen examples. The Siamese architecture means parameter sharing for processing both inputs, such that all embeddings are calculated by the same set of parameters, see Figure 4. For our base

[1] we write $(\cdot)_+$ for $\max\{0, \cdot\}$

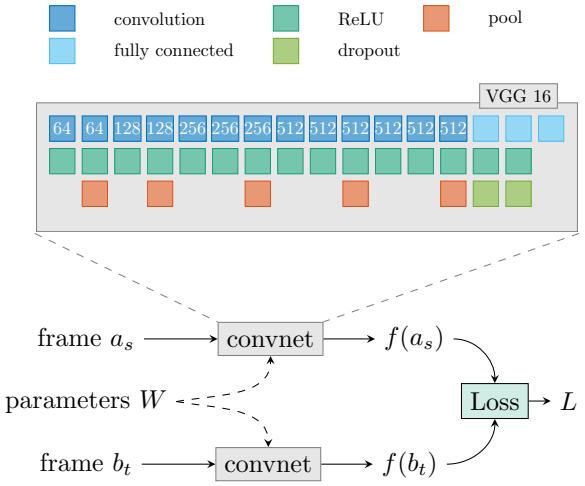

Figure 4: *Each frame x of the pair (a_s, b_t) is encoded by our convnet as two 1024 dimensional vectors $f(a_s)$ and $f(b_t)$. Training the network parameters θ is done by minimizing the contrastive loss from Eq.* (1)

convnet model, we choose the VGG-16 neural network layout [Simonyan and Zisserman 2014], for it has demonstrated state-of-the-art performance in various computer vision tasks. We then adapt the model by introducing an additional fully connected layer with 1024 output neurons. We fine-tune only the last five layers of the VGG and the additional embedding layer to adapt to the intended similarity metric while keeping the first few layers frozen.

3.4 Finding the Optimal Matching Tour

The second part of our approach is dedicated to finding a matching tour through the similarity matrices, i.e., rejecting implausible frame matches even when they are predicted as similar by our neural network. This step essentially introduces the time dimension into the algorithm, yielding a full synchronization of the processed videos. It also acts as a filtering algorithm for the data, removing mislabelled pairs and thus providing more reliable training data for the next interpretation phase. From the current state of the trained neural network we are given a metric δ and a cost matrix C as described above. We then pre-process the cost-matrix and apply a path-finding step.

Preprocessing Particularly in the early iterations, the similarity matrices produced by our neural network contain a lot of noise, since it is not yet fully trained for the task. This noise exhibits a low-rank structure, since any frame that is not correctly embedded is likely to corrupt an entire column (or row) of the similarity matrix, see top-row of Figure 5. Additionally, some of the pairs are, indeed, rather similar although we would actually like to treat them as different. For example, many journeys through rural areas with little information but crop fields on both sides of the road appear extremely similar.

Observing this kind of regularity in the noise allows for a filtering using a low-rank approximation. This is achieved by using the singular value decomposition $U\Sigma V^\star$ of the cost matrix C which gives $C_n = \sum_{k=1}^{n} U_k \Sigma_k V_k^\star$ as the n-th low-rank approximation. We refer to $C - C_n$ as the "cleaned" cost matrix. The result is illustrated in the mid-row of Figure 5 using $n = 20$. In the following we only refer to the cleaned version of the cost matrix. Abusing notation we would further denote the cleaned cost matrix as C.

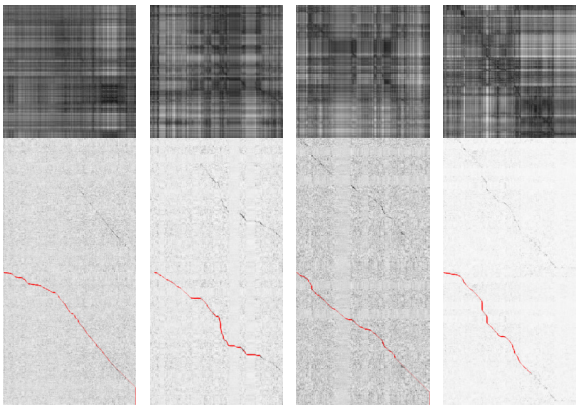

Figure 5: *The upper row shows the computed similarity matrices from the first iteration in our training process. As they are clearly noisy, directly applying Dijkstra's algorithm would fail. The lower row shows the cleaned cost matrices based on the observation that the noise can be approximated by a low-rank matrix. This preprocessing makes finding a path significantly robuster. In all the cleaned matrices, the correct path was found.*

Finding a matching tour After getting the cost matrices from our neural network and filtering out noise in the preprocessing phase, the only step missing for aligning a video pair is to find a plausible path through its respective cost matrix. We formulate this problem as a shortest-path problem. Hereby, each frame pair (a_i, b_j) included in the path penalizes the length of the entire tour by c_{ij}. To remove any human interaction we tackle the problem of video synchronization under uncertainty regarding the location of the partial overlap of a given video pair in contrast to [Wang et al. 2014]. There might even be a road-detour, forcing a correct alignment to include separate parts of the videos. Non-matching regions might corrupt the entire path when searching for the global optimum over the entire matrix. Therefore, we partition the cost matrix along one time dimension into multiple column-stripes. The stripes are 35 pixels wide corresponding to 70 frames, a parameter which was determined from experiments as a good point in the middle between having too much and too little temporal information.

In order to ensure that paths are not bound to start at the top-left corner of each stripe, we further add zero-padding to the left edge of the stripes, i.e. the time dimension of only one of the videos (in contrast to [Wang et al. 2014]). This enables almost complete freedom regarding the location of each match and limits cases of stripes suffering from imprecise information introduced by non-matching stripes. Following the matching of all stripes, we then stitch all those local matches back together to get a robust global matching tour which does not necessarily span through the entire cost matrix.

Finding a local matching tour in each stripe is done by applying Dijkstra's algorithm from S_i to S_{i+1} (see Figure 7). We only allow three directions, downwards, rightwards and a diagonal bottom-right step, which prevents circular tours and drastically reduces the computational complexity. Each possible path

$$P = (S_i, p_1, p_2, \ldots, p_n, S_{i+1}) \in \mathcal{P} \tag{2}$$

from S_i to S_{i+1} has an associated *matching cost* defined as

$$\pi(P) = \sum_{i=1}^{n} w(c_{p_i}). \tag{3}$$

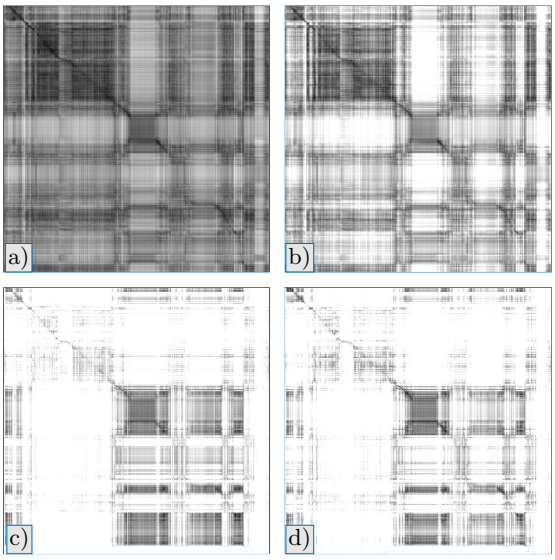

Figure 6: *Figures a) through d) show the cost matrix of the same video pair throughout the iterative training iterations. Note how the path becomes clearer and more distinct.*

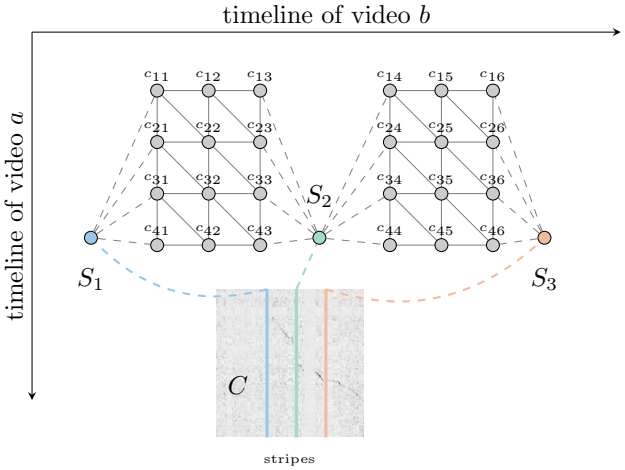

Figure 7: *Given a cost matrix C, we split up one video by dividing the cost matrix into stripes. Each entry in C represents the similarity of the corresponding frame-pair. Finding the local shortest path in each stripe independently and testing the global plausibility by stitching them together, we get reliable matching tour through for synchronizing video a and b.*

We use $w(x) := x^6$ to add a higher penalty for bad frame-matches. Considering each stripe individually might generate locally optimal paths, which are not necessarily a part of the correct global matching tour. So instead of combining all local optimal paths, we compute a set of candidate paths $\mathcal{P}'_S$ from all possible paths $\mathcal{P}_S$ per stripe, which meets the following properties:

1. The bound $\pi(P') < \eta \cdot \pi(P)$ holds for $P' \in \mathcal{P}'_S$ and random paths $P \in \mathcal{P}_S$. A true path segment should have significantly lower matching cost than a random path in the current stripe. We use $\eta = 0.215$.
2. The path has non-zero costs, i.e. $\pi(P') \neq 0$ for $P' \in \mathcal{P}'_S$ to avoid trivial self-alignment from noisy tagged video duplicates.
3. The path $P' \in \mathcal{P}'_S$ does not feature a longer region of constant horizontal or vertical slop. This normally occurs when one video is longer than the other and a zero-padding is added to keep the cost matrix square.
4. The path $P' \in \mathcal{P}'_S$ can be seamlessly stitches together with a candidate path P'_-/P'_+ from the previous/following stripe with a tolerance of τ frames in each direction.

From all remaining paths the path with minimal costs is returned. Parameters of this heuristic are more conservatively applied in the early stages of training since the not yet fully optimized neural network produces a lot of noise. These are later relaxed, mostly in the form of relaxing the thresholds mentioned above. In practise, this heuristics are implemented as an indicator function $\Gamma(P)$, which adds ∞ to the matching costs in (3) if one of these constraints is violated.

3.5 Generating the Training Data

After calculating the cost matrices and extracting the matching tours, we can use the extracted data for further training of the convolutional neural network. The quality of the training data has an extensive impact on the success of any machine learning algorithm. Too simple training samples (like consecutive frames for positive pairs) would allow the neural network to "cheat", i.e. learning Euclidean distance of the frame pixels $\delta'(a_i, b_j)$ without good generalization. On the other hand, in order to generate highly diverse training samples, one would either need a fully trained neural network or labelled data. To resolve this chicken-egg problem we iteratively add robustly classified training data, exploiting the current state of the distance measure combined with the robust tour estimation. By doing so, we also apply the concept of curriculum learning proposed in [Bengio et al. 2009] which showed that training a neural network with an increasing data complexity helps to achieve a quicker convergence.

While random sampling of negative training pairs is sufficient in the initial training iteration, these quickly become dully simple. Therefore, we mix them with more challenging training examples obtained frame pairs which are part of local candidate paths $\mathcal{P}'_S$ but not used in the global matching tour. Initially, we had expected that using only those challenging frame-pairs would result in a strong, useful set of negative training samples. In practice however, picking only hard samples had a negative influence on the training process, as many of the pairs were challenging even for human observers, like straight roads along crop fields with no further information.

Initial Training Data Each video file in the dataset is only tagged with a date. To initially generate training pairs, we use consecutive frames within arbitrary videos, i.e. extracting frames $(a_t, a_{t+\eta})$ with η sampled uniformly in $\eta \sim \mathcal{U}([1, x])$, with x initialized to 5 and slowly increased to 25 with convergence (this also supports our curriculum leaning scheme). Negative examples are simply random frames (a_t, a_s) with a minimal margin between t and s of 500 frames. Although we could have randomly sampled frames from different videos here, this would have an additional variation in image content, i.e. weather conditions, which would have taught the network to rely on vague visual cues.

Re-Generate Training Data Assuming the first stage of training using the initial training data has converged, we randomly select video pairs $a, b \in \mathcal{V}, a \neq b$ with a distance of at most d days. We

start by embedding the dataset and calculating the similarity matrices for all video pairs within a date range d of 10 days. While we could actually compare all pairs regardless of date differences, we have noticed a drop in both the quality and the quantity of the results for larger values of d. This could be explained by the nature of the videos, demonstrating lesser difference within a short period of time comparing to videos taken months apart. Furthermore the variation in daylight, road conditions and weather is likely to remain rather consistent over a short period of time.

We then look for all matching tours in the analysed data as outlined in Section 3.4 and used our freshly labelled data for a further optimization for the convnet. We repeat these steps while increasing d with each iteration first to 30 and then to 60. This is accompanied by an increase in the amount of matches even within the initial 10 day range. The number of correctly recognized pairs rose from $55,355$ pairs after the first iteration to $207,890$ after the second. The 60 day range seemed to suffice to link the different seasons. Additionally, the increased diversity in the training data leads to higher accuracy and reduced noise at the embedding-phase. This can be seen in Figure 6 where the exact same pair was compared during the different stages of training.

Note that the independence from human intervention as well as the lack of labelled data during the training phase could eventually lead to some generated training data, especially in early iterations, being mislabelled. However, we have witnessed this to only affect a very small proportion of the data (roughly 2%, changes between iterations) and due to the large amounts of data we use to train our convnet, this exhibits a little to no impact on our results.

4 Applications

The presented learning of visual similarity in scenes yields a dimensionality reduction which maps full HD video frames to 1024 dimensional embedding vectors. This reduces the total size of the video collection from 3.5 TB to 181 GB. Instead of touching videos directly their embeddings can be used instead which enables new applications in computer graphics. In addition to video synchronisation we outline two additional applications, namely querying an entire snippet and the extension of composing a video across several seasons.

4.1 Video Query With Consistency over Time

One application of this neural network encoding allows for a very efficient video similarity search, querying for all related video snippets with the same route/location. Hence, it avoids the need for an expensive nearest neighbour search based on multiple SIFT features per frame or alternative encoding techniques like VLAD [Jegou et al. 2010] and Fisher-Vectors [Perronnin and Dance 2007].

Finding all similar frames is usually be done by applying an exhaustive search over all possible embeddings. Rather than finding only similar frames, we are interested in extracting all frames from which adjacent frames are also reasonable close *over time* wrt. the query sequence. It seems to be reasonable to include these additional information from the surrounding frames of the query for robust *time-window-based* queries.

Therefore, we propose to combine the learned per-frame similarity metric δ and the globally optimal path extraction from the cost matrix to allow these time-window-based queries. Let $q_s \in \mathcal{V}$ be a query frame for some $s \in \mathbb{N}$. And let $k \in \mathbb{N}$ be a time window, such that the video sequence consists of frames

$$Q = (q_\ell)_{\ell=s-k,s-k+1,\dots,s+k}, \quad q_\ell \in \mathcal{V}.$$

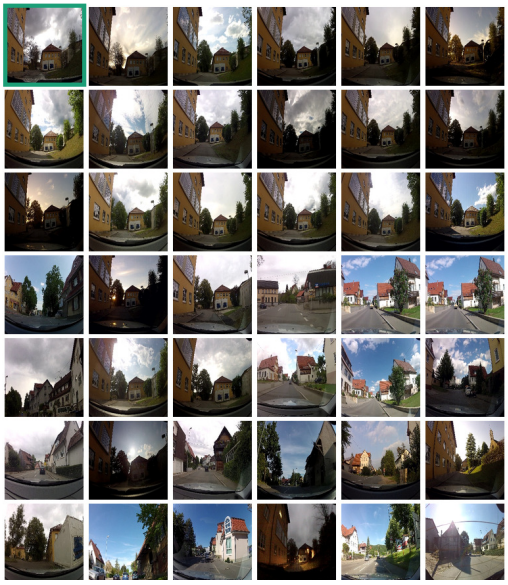

Figure 8: *Querying a sequence (left top corner) in our entire collection of 602 video gives several candidate frames ranked according the costs of found matching tour (from top left to bottom right). The time-window is 140 frames.*

Ranking the cost matrices C_i between Q and all N video snippets $v_{i=1,2,\dots,N}$ according the costs $(c_i)_{i=1,\dots,N}$ of the matching tour gives a list of possible candidates of related video snippets (see Figure 8).

Although this video retrieval approach might be significantly accelerated by some sort of advanced index structure, it processes 140 key-frames per second.

4.2 Compositing Videos with automatic transitions

Interpreting the output of the neural network as matching similarities, one might use these information to find optimal cut transitions between multiple video takes. Not only could the cost matrices be used for aligning more than one video pair (a, b) at the same time – one could also globally optimize a particular key frame in order to transition from video a to video b.

First, for a given frame we query all related video snippets with partial overlap. Then, we use the resulting cost matrices to find the optimal transitions.

5 Experiments and Results

We have conducted several experiments to asses the performance of our algorithm. The following results are obtained on a workstation with an Nvidia Titan X GPU. We demonstrate the robustness on aligning a couple of challenging scenes. To evaluate our method despite missing ground truth, we pursue two approaches. First, we compared our convnet prediction of similarity to the SIFT-based histogram matching from [Wang et al. 2014]. Next, we add an artificial blur to the input frames (though they already contain motion blur) or blank out an area in the image to examine the robustness of both approaches.

5.1 Performance

The VGG convnet requires the resizing of input data to 227×227 pixels. This necessity adds a small overhead to the processing time, which could be saved by executing as a pre-processing step, once for the entire dataset. The large quantity of data makes it impossible to pre-load the entire dataset to the GPU or to the RAM. Yet we have witnessed a significant decrease in loading time (IO time) when, prior to training/embedding, storing the pre-processed video data on a local SSD. When loading the raw data from the file server, the system is able to align 70 frames per second including pre-processing, embedding, generation and clean-up of the similarity matrix and path optimization. During training with a batch size of 50 pairs and the pre-processed data loaded from the file sever, we have experienced an IO time of $\sim$ 17 seconds and $\sim$ 1.7 seconds for the forwards and backwards computations on the GPU.

5.2 Robust Alignment

In order to demonstrate the robustness of the whole alignment procedure, videos in various different settings are aligned. Figure 1 shows the successful alignment through different weather conditions with and without headlights. Disturbances on the windscreen such as dirt (Figure 9a), raindrops or a sweeping wiper (Figure 9b) hardly influence the estimated path. The approach is further robust against varying traffic conditions, parked cars and even construction sites in various stages. Various camera orientations and even drastic changes in the camera position, e.g. due to a lane change (up to 2 meters) are correctly handled, see Figure 10.

5.3 Comparison to a SIFT Histogram Based Approach

We have re-implemented the approach of extracting SIFT features from unprocessed video frames (approx. 1k SIFT features each) in full HD resolution and applying a fast approximate of a nearest neighbour search. Please note, that finding SIFT features was applied to the high-res image (1920×1080 pixels), while our network operates on a down-scales image (227×227 pixels). Scaling down the frames would most certainly result in a more reasonable running time, yet, as can be seen in Figure 11, even for the full size images, the SIFT based method fails to find enough common features to properly align two videos.

Qualitative Performance For a fair evaluation we add the step of cost cleaning, described in subsection 3.4, based on the low-rank approximation of the cost matrix to the pipeline of [Wang et al. 2014] with a manual tuning of the rank approximation. Still, we were unable to use the SIFT histogram approach to align videos of acquisition dates exceeding 4 months, see Figure 11. SIFT fails to match corresponding frames from these videos probably due to a wide variance in appearance, e.g. a snow covered surrounding compared to a captured landscape in spring.

Runtime The usage of a neural network compares favourably to a SIFT based approach concerning the run time and storage requirements. With a relative small footprint of 550MB to store the weights it can process 70 frame-pairs per second, compared to the gigabytes and 9 fps required for matching SIFT vectors using a state-of-the-art GPU based nearest neighbour search[2].

[2]The library FLANN [Muja and Lowe 2009] required more than 4 days.

5.4 Training Neural Network

To solve the video alignment problem, we used a pre-trained VGG in our Siamese Neural-Network. Training and evaluation experiments were done using the python framework *Thenao* [Bastien et al. 2012]. The optimization method used was RMSprop [Tieleman and Hinton 2012] in mini-batches of 50 frame-pairs.

The contrastive loss hyper-parameter m is initialized to the median radius calculated from the pairwise distances over 10 mini-batches. This helps keeping the gradients from getting too large or too small for the negative pairs and resembles the pseudo-hard negatives mentioned by [Schroff et al. 2015].

We observed no convergence when training solely with hard negative training samples, i.e. frame-pairs for which the neural network thinks that they are similar but they do not pass our heuristic for the global matching path. Therefore, our negative dataset was compiled with an equal amount of easy, random samples and those hard samples.

5.5 Testing Robustness of the Video Retrieval

Instead of computing SIFT matching-histogram, the neural network is able to encode relationships within scenes, leading to better performance through understanding the context. The application of the Dijkstra-based matching algorithm allows for coherent matches over multiple frames. To challenge our approach we manipulate *one* video of the video pair in two ways:

Corrupt Video Parts For exploring the robustness of this method beyond videos of dusk, rain, fog, night, we applied a Gaussian blur to the full-sized frame pairs. Whereas SIFT-based synchronization failed to find a proper match, it made no difference to the convnet.

Blank Frame Areas Another way of testing robustness was to replace a part of a frame by a black rectangle.

In both of the above mentioned cases, the proposed method had no trouble finding the correct alignment tour. SIFT-based matching, however, started failing even on small disturbances. See the supplementary material for the respective results.

6 Limits and possible Future Work

Our technique is subjected to a number of limitations. First, ignoring the high-res input drops discriminative local information and probably results in less accurate estimation of the similarity function. This could be optimized by experimenting with different neural network architectures such as the recent residual networks [He et al. 2015]. Since the local information plays an important role in the accuracy of the system, it could benefit from integrating the dimensionality reduction, currently done in pre-processing, into the network in the form of larger pooling kernels. However, this would render the usage of pre-trained weights though impossible and would require a sufficient, high resolution dataset to pre-train on.

Another modification to the network could be cosine similarity loss function instead of the contrastive loss. As recently shown in [Cheng and Kartsaklis 2015], the cosine similarity tends to improve accuracy comparing to the contrastive loss in some NLP tasks. We would therefore like to explore its advantages when applied to vision related tasks.

At the moment, we apply a tough constraint to the Dijkstra implementation by only allowing it progress in three directions. This has

Figure 9: *Some pairs of synchronized Videos across different months and different weather situations. Notably the algorithm can robustly handle windscreen wipers and large time differences.*

Figure 10: *Various correctly matched frame-pairs with drastic changes in scene content, e.g. weather, sky appearance, illumination, traffic and road conditions, road signs and construction work.*

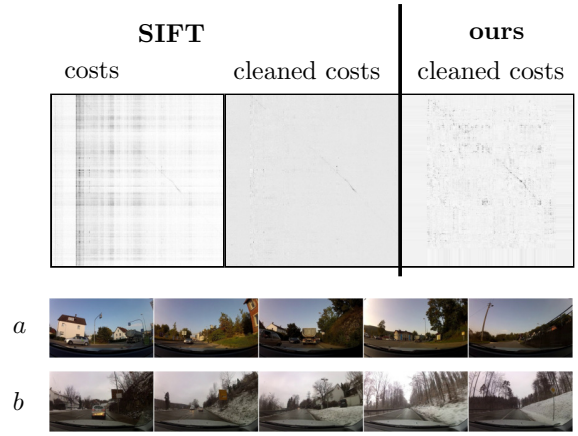

Figure 11: *Although, there should be a matching tour through the entire video snippet, it is not possible to align both videos using SIFT features [Wang et al. 2014]. Rows a and b show a few frames from each video taken in September 2012 and February 2013 respectively. Training from frame-pairs with a large time horizon allows to learn similarities even between these videos. The cost matrix in our approach contains reasonable information for most frames.*

proven sufficient for non-linear matches and we have not yet encountered videos where it has lead to false negatives. Yet, as the network becomes more robust, we hope to no longer require this kind of constraint.

7 Conclusion

To our knowledge, we present the first application of a deep convolutional neural network to align videos without any human intervention and no need for manually created labels. During training, the system autonomously harvests training samples from a large video dataset of driving sequences, alternating between improving the learned similarity measure and robustly aligning videos for the next iteration. Alignment is further improved by removing the low-rank noise from computed similarity matrices and by exploiting Dijkstra's algorithm to estimate a globally optimal path. We demonstrate that the proposed approach manages to cope well with videos showing drastic changes in environment settings caused by road conditions or seasonal effects. Our method tolerates large changes in appearance that cannot be handled by other existing methods. Its practicability is demonstrated in several applications namely video synchronization, video query and compositing of seasonal-varying videos.

This neural network as a black-box in the fashion of transfer learning to predict similarity between videos frames could also be embedded in other video applications. Finding similar frames in massive video collections using frame embeddings and matching in small window frames allows for finding matches which are also plausible over time. Describing the entire scene by one embedding vector rather than a collection of local features further reduces the computational complexity of exhaustive nearest neighbour search. The described training procedure can be transferred to other methods of video manipulation.

References

AGARWALA, A., ZHENG, K. C., PAL, C., AGRAWALA, M., COHEN, M., CURLESS, B., SALESIN, D., AND SZELISKI, R. 2005. Panoramic video textures. In *ACM Transactions on Graphics (SIGGRAPH)*, ACM, New York, NY, USA, SIGGRAPH '05, 821–827.

BASTIEN, F., LAMBLIN, P., PASCANU, R., BERGSTRA, J., GOODFELLOW, I., BERGERON, A., BOUCHARD, N., WARDE-FARLEY, D., AND BENGIO, Y. 2012. Theano: new features and speed improvements. *arXiv preprint arXiv:1211.5590*.

BENGIO, Y., LOURADOUR, J., COLLOBERT, R., AND WESTON, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, ACM, 41–48.

BROMLEY, J., GUYON, I., LECUN, Y., SCKINGER, E., AND SHAH, R. 1994. Signature verification using a "siamese" time delay neural network. In *In NIPS Proc*.

BRYAN, N., SMARAGDIS, P., AND MYSORE, G. 2012. Clustering and synchronizing multi-camera video via landmark cross-correlation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2389–2392.

CASPI, Y., AND IRANI, M. 2002. Spatio-temporal alignment of sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence 24*, 1409–1424.

CHENG, J., AND KARTSAKLIS, D. 2015. Syntax-aware multi-sense word embeddings for deep compositional models of meaning. *arXiv preprint arXiv:1508.02354*.

HADSELL, R., CHOPRA, S., AND LECUN, Y. 2006. Dimensionality reduction by learning an invariant mapping. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, vol. 2, 1735–1742.

HE, K., ZHANG, X., REN, S., AND SUN, J. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.

JEGOU, H., DOUZE, M., SCHMID, C., AND PREZ, P. 2010. Aggregating local descriptors into a compact image representation. In *CVPR*, IEEE Computer Society, 3304–3311.

KANG, S. B., UYTTENDAELE, M., WINDER, S., AND SZELISKI, R. 2003. High dynamic range video. *ACM Trans. Graph. 22*, 3 (July), 319–325.

LECUN, Y., JACKEL, L., BOTTOU, L., CORTES, C., DENKER, J. S., DRUCKER, H., GUYON, I., MULLER, U., SACKINGER, E., SIMARD, P., ET AL. 1995. Learning algorithms for classification: A comparison on handwritten digit recognition. *Neural networks: the statistical mechanics perspective 261*, 276.

LEI, C., AND YANG, Y.-H. 2006. Tri-focal tensor-based multiple video synchronization with subframe optimization. *IEEE Transactions on Image Processing 15*, 9, 2473–2480.

LIU, C., YUEN, J., AND TORRALBA, A. 2011. Sift flow: Dense correspondence across scenes and its applications. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 33*, 5 (May), 978–994.

MOBAHI, H., COLLOBERT, R., AND WESTON, J. 2009. Deep learning from temporal coherence in video. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ACM, 737–744.

MUJA, M., AND LOWE, D. G. 2009. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application VISSAPP'09)*, INSTICC Press, 331–340.

NGO, C.-W., MA, Y.-F., AND ZHANG, H.-J. 2005. Video summarization and scene detection by graph modeling. *Circuits and Systems for Video Technology, IEEE Transactions on 15*, 2 (Feb), 296–305.

PERRONNIN, F., AND DANCE, C. R. 2007. Fisher kernels on visual vocabularies for image categorization. In *CVPR*, IEEE Computer Society.

RÜEGG, J., WANG, O., SMOLIC, A., AND GROSS, M. 2013. Ducttake: Spatiotemporal video compositing. In *Computer Graphics Forum*, vol. 32, Wiley Online Library, 51–61.

SAND, P., AND TELLER, S. 2004. Video matching. *ACM Trans. Graph. 23*, 3 (Aug.), 592–599.

SCHARSTEIN, D., AND SZELISKI, R. 2002. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision 47*, 1 (May), 742.

SCHROFF, F., KALENICHENKO, D., AND PHILBIN, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.

SHRSTHA, P., BARBIERI, M., AND WEDA, H. 2007. Synchronization of multi-camera video recordings based on audio. In *Proceedings of the ACM International Conference on Multimedia (ACMMM)*, ACM, New York, NY, USA, MM '07, 545–548.

SIMONYAN, K., AND ZISSERMAN, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR abs/1409.1556*.

TIELEMAN, T., AND HINTON, G., 2012. Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural Networks for Machine Learning.

TUYTELAARS, T., AND GOOL, L. J. V. 2004. Synchronizing video sequences. In *CVPR (1)*, 762–768.

UKRAINITZ, Y., AND IRANI, M. 2006. Aligning sequences and actions by maximizing space-time correlations. In *Proceedings of the 9th European Conference on Computer Vision - Volume Part III*, Springer-Verlag, Berlin, Heidelberg, ECCV'06, 538–550.

WANG, O., SCHROERS, C., ZIMMER, H., GROSS, M., AND SORKINE-HORNUNG, A. 2014. Videosnapping: Interactive synchronization of multiple videos. *ACM Trans. Graph. 33*, 4 (July), 77:1–77:10.

WIESCHOLLEK, P., WANG, O., SORKINE-HORNUNG, A., AND LENSCH, H. P. 2016. Efficient large-scale approximate nearest neighbor search on the gpu. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

WOLF, L., AND ZOMET, A. 2006. Wide baseline matching between unsynchronized video sequences. *International Journal of Computer Vision 68*, 1, 43–52.

ZHOU, F., AND DE LA TORRE FRADE, F. 2009. Canonical time warping for alignment of human behavior. In *Advances in Neural Information Processing Systems Conference (NIPS)*.